

Random Geometry on the Sphere

Jean-François Le Gall*

Abstract. We introduce and study a universal model of random geometry in two dimensions. To this end, we start from a discrete graph drawn on the sphere, which is chosen uniformly at random in a certain class of graphs with a given size n , for instance the class of all triangulations of the sphere with n faces. We equip the vertex set of the graph with the usual graph distance rescaled by the factor $n^{-1/4}$. We then prove that the resulting random metric space converges in distribution as $n \rightarrow \infty$, in the Gromov-Hausdorff sense, toward a limiting random compact metric space called the Brownian map, which is universal in the sense that it does not depend on the class of graphs chosen initially. The Brownian map is homeomorphic to the sphere, but its Hausdorff dimension is equal to 4. We obtain detailed information about the structure of geodesics in the Brownian map. We also present the infinite-volume variant of the Brownian map called the Brownian plane, which arises as the scaling limit of the uniform infinite planar quadrangulation. Finally, we discuss certain open problems. This study is motivated in part by the use of random geometry in the physical theory of two-dimensional quantum gravity.

Mathematics Subject Classification (2010). 05C80, 60D05

Keywords. Planar map, Gromov-Hausdorff convergence, Brownian map.

1. Introduction

In the last ten years, there has been much interest in discrete and continuous models of random geometry in two dimensions. As a first naive attempt to construct a continuous model, one could imagine choosing at random a Riemannian metric on the sphere. However there seems to be no canonical way of making such a choice, and a better approach leading to a continuous object that is universal in some sense is to start from discrete models. Informally, we will consider a large graph drawn on the sphere and chosen at random in a suitable class: We then expect that the suitably rescaled graph distance on the vertex set will converge in an appropriate sense, when the size of the graph tends to infinity, to a random metric on the sphere.

More precisely, we consider planar maps, which are finite connected graphs embedded in the sphere \mathbb{S}^2 – more precise definitions will be given in Section 2 below. One is interested in the “shape” of the graph, so that two planar maps are identified if the second one is the image of the first one under an orientation-preserving homeomorphism of the sphere. The faces of the map are the connected

*This paper will appear in the Proceedings of ICM Seoul 2014

components of the complement of edges, and a planar map is called a triangulation if all faces are triangles, possibly with two edges glued together. For technical reasons, it is convenient to deal with *rooted* triangulations, meaning that there is a distinguished oriented edge called the root edge, whose tail is the root vertex. Thanks to the preceding identification, for every fixed even integer n , there are only a finite number of rooted triangulations with n faces, and therefore it makes sense to choose one of them at random, which we denote by T_n . Then, supposing that there is a canonical way of embedding T_n in the sphere (recall that the embedding of a planar map is only defined up to orientation-preserving homeomorphisms), it seems reasonable to conjecture that the vertex set $V(T_n)$ of T_n will become dense in \mathbb{S}^2 when $n \rightarrow \infty$, and that the graph distance d_{gr} on this vertex set will converge, modulo a suitable rescaling, to a random metric on \mathbb{S}^2 .

The preceding program is still the subject of active research – there are indeed (almost) canonical ways of embedding triangulations, using circle packings or the uniformization theorem for Riemann surfaces, see the discussion in Section 6 below. Here we will adopt a slightly different point of view. We consider the finite metric space $(V(T_n), d_{\text{gr}})$ as a (random) element of the set \mathbb{K} of all isometry classes of compact metric spaces. The space \mathbb{K} is equipped with the Gromov-Hausdorff distance d_{GH} (cf. subsection 3.1), and (\mathbb{K}, d_{GH}) is both separable and complete. A key result [31], which solves a conjecture of Oded Schramm [44], states that

$$(V(T_n), 6^{1/4} n^{-1/4} d_{\text{gr}}) \xrightarrow[n \rightarrow \infty]{(d)} (\mathbf{m}_\infty, D), \quad (1)$$

where the convergence holds in distribution in the space (\mathbb{K}, d_{GH}) , and the limit (\mathbf{m}_∞, D) is a random metric space, which is called the Brownian map after Markert and Mokkadem [37]. Despite the somewhat unusual fact that we are dealing with random compact metric spaces, (1) is nothing but a particular case of the standard notion of convergence in distribution for random variables with values in a Polish space. Note that the constant $6^{1/4}$ in (1) is just a normalization factor.

A very important feature of the convergence (1) is the fact that it holds for much more general random planar maps, with the *same* limiting space (\mathbf{m}_∞, D) up to unimportant scaling constants. Recall that a planar map is a p -angulation if all faces have degree p . Then it was proved in [31] that an analog of (1) holds for p -angulations with a fixed number of faces, for every even integer $p \geq 4$. Note that the special case of quadrangulations, corresponding to $p = 4$, has been treated independently by Miermont [41] via a different method. Similarly, analogs of (1) hold for general planar maps with a fixed number of edges [11] or for bipartite planar maps with a fixed number of edges [1]. It is also possible to impose local constraints on the planar maps: A result similar to (1) holds for simple triangulations or quadrangulations [2], where no loops or multiple edges are allowed, or for quadrangulations with no pendant vertices [8]. One indeed expects that the Brownian map will appear as the scaling limit of very general random planar maps provided some bound holds on the distribution of the degree of a typical face (on the other hand, the paper [33] shows that different scaling limits may occur if one considers distributions that favour the appearance of “very large” faces).

The Brownian map thus appears as a *universal* object, in the sense that it is the scaling limit of many different models of random planar maps: This is of course similar to the case of Brownian motion, which is the universal scaling limit of many different random walks on the lattice. Just as Brownian motion can be viewed as a purely random continuous curve, the Brownian map seems to be the right model for a purely random surface. Note that the Brownian map (\mathbf{m}_∞, D) is almost surely homeomorphic to the sphere \mathbb{S}^2 [35] (see also [39]), even though its Hausdorff dimension is equal to 4. However, there is no canonical homeomorphism and we cannot a priori use the Brownian map to obtain a “canonical” random metric on the sphere as suggested by preceding remarks.

The main goal of the present work is to present a general theorem of convergence toward the Brownian map (Theorem 3.1 below), which was derived in the series of papers [29, 30, 31]. We also give a detailed construction of the Brownian map, by showing that this random compact metric space can be obtained by gluing certain pairs of points in another famous probabilistic model, the Brownian continuum random tree or CRT, which was introduced and studied by Aldous [3, 4]. Note that the CRT itself is a universal scaling limit of discrete trees, in a sense analogous to (1). The best way to understand the construction of the Brownian map from the CRT is to start from the discrete bijections that exist between various classes of planar maps and corresponding classes of labeled trees. For this reason, we start in Section 2 by explaining these bijections in the particular case of quadrangulations (there exist similar bijections for triangulations, or more generally for p -angulations, but their description is more complicated). In Section 3, after stating the main theorem of convergence that extends (1), we give a precise definition of the CRT as the tree coded by a normalized Brownian excursion, and we construct the Brownian map via the above-mentioned gluing procedure. Note that, although the discrete bijections between planar maps and trees are very far from being sufficient to get results such as (1), they provide useful insight into the construction of the Brownian map. In Section 4, we describe the detailed results that are known about geodesics in the Brownian map. Section 5 is devoted to the Brownian plane, which can be viewed as an infinite-volume version of the Brownian map. The Brownian plane shares many properties of the Brownian map, but it enjoys an additional scale invariance property, which makes it more suitable for explicit calculations (we briefly present some recent results from [22]). One can also view the Brownian plane as a continuous analog of the infinite random lattices known as the uniform infinite planar triangulation and quadrangulation, which have been studied extensively in the recent years. Finally, Section 6 discusses what is probably the most important open problem in the area, namely finding a canonical construction of the Brownian map in terms of a random metric on the sphere. A related discussion can be found in Benjamini [9].

To complete this introduction, let us mention that, although our main motivation for the following results came from probability theory, there are important connections with several other areas of mathematics and physics. In particular, from the point of view of combinatorics, one can derive information about properties of “typical” large planar maps from the known results about the Brownian

map. For instance, the fact that the Brownian map is homeomorphic to the sphere implies that for a large planar map chosen uniformly at random in a given class (e.g. a large triangulation) there cannot exist a cycle whose length is small in comparison with the diameter of the graph, such that both regions separated by this cycle have a macroscopic size (see [35]). Similarly, the statements about geodesics in the Brownian map show that, in a typical large planar map, the geodesic between two vertices chosen at random is “macroscopically unique”, meaning that the distance between two geodesics will be small in comparison with the diameter of the graph. Another strong motivation for this work came from physics, and particularly from the use of large random planar maps as models of random geometry in two-dimensional quantum gravity (see the book [5]). In this connection, we mention the important work of Bouttier and Guitter (see in particular [13, 14, 15]) which motivated part of our results. It is worth mentioning that a different mathematical approach to two-dimensional quantum gravity relying on the Gaussian free field has been given by Duplantier and Sheffield [24]. The construction of [24] seems quite different from our perspective, but the paper [45] gives a number of conjectures that would relate large random planar maps and the Brownian map to the Gaussian free field approach. The very recent work of Miller and Sheffield [42] seems promising in this respect.

Acknowledgements. I thank Nicolas Curien and Igor Kortchemski for their help with figures.

2. Planar maps and bijections with trees

In this section, we recall the basic definitions concerning planar maps and we explain how planar maps can be coded by trees, in the particular case of quadrangulations. This coding is an important ingredient of the proofs, and it also helps to understand the definition of the Brownian map that will be given below.

2.1. Planar maps. Let us start with the definition of a planar map.

Definition 2.1. *A planar map is a proper (without edge-crossing) embedding of a finite connected graph in the two-dimensional sphere \mathbb{S}^2 . A rooted planar map is a planar map given with a distinguished oriented edge, which is called the root edge and whose tail vertex is called the root vertex. Two rooted planar maps are identified if they correspond via an orientation-preserving homeomorphism of \mathbb{S}^2 .*

We in fact allow loops and multiple edges in our graphs. Following the classical terminology found in combinatorics, the word “graph” should be replaced by “multigraph” in the preceding definition. The faces of a planar map are the connected components of the complement of the union of edges, and the degree of a face counts the number of edge sides that are incident to this face (in particular it may happen that both sides of an edge are incident to the same face, and then this edge is counted twice in the degree of the face).




Figure 1. A rooted quadrangulation with 7 faces

For any integer $p \geq 3$, a planar map is called a p -angulation (a triangulation if $p = 3$, a quadrangulation if $p = 4$) if all its faces have degree p . Fig. 1 shows an example of a quadrangulation with 7 faces. For every integer $n \geq 1$, we write \mathcal{M}_n^p for the set of all rooted p -angulations with n faces. When p is odd, the set \mathcal{M}_n^p is empty if (and only if) n is odd. So when we deal with odd values of p (in particular with triangulations), we will always assume that n is even.

Thanks to the identification in Definition 2.1, the sets \mathcal{M}_n^p are finite. Enumeration results for these sets were obtained by Tutte (see in particular [46]) in a series of important papers motivated by the four-color theorem.

We will use the notation d_{gr}^M for the graph distance on the vertex set $V(M)$ of a planar map M .

2.2. Bijections with trees. In this section we explain the Cori-Vauquelin-Schaeffer bijection between rooted quadrangulations and well-labeled trees [18, 19]. We restrict ourselves to the special case of quadrangulations for the sake of simplicity, but similar bijections exist for p -angulations for any p (see in particular [12]).

First recall that a plane tree τ is a (finite) rooted ordered tree: A way of specifying a plane tree is to represent each of its vertices by a finite word made of positive integers, in such a way that the empty word \emptyset corresponds to the root or ancestor of the tree, and that, for instance, the word 13 corresponds to the third child of the first child of the root (see the left side of Fig. 2).

Then a well-labeled tree is a plane tree τ , with vertex set $V(\tau)$, whose vertices v are assigned labels $(\ell_v)_{v \in V(\tau)}$, in such a way that the following properties hold. First, the label of the root is 1, then the label ℓ_v of any vertex v is a positive integer, and finally $|\ell_v - \ell_{v'}| \leq 1$ if v and v' are adjacent vertices.

With any well-labeled tree $(\tau, (\ell_v)_{v \in V(\tau)})$ with n edges, we can associate a rooted quadrangulation Q with n faces via the following construction. Suppose that the tree is embedded in the plane in the (obvious) way as suggested in the




Figure 2. The Cori-Vauquelin-Schaeffer bijection. On the left side, a well-labeled tree (the framed numbers are the labels assigned to the vertices). On the right side, the edges of the associated quadrangulation Q appear in thick curves.

left part of Fig. 2 (in particular the successive children of a vertex appear from left to right). Then the vertex set of Q will be the union of the vertex set of τ and of an extra vertex, and we now explain how to construct the edges of Q . First recall that a corner of the tree τ is an angular sector between two successive edges of τ around a given vertex. The set of all corners of τ is given a cyclic ordering by moving clockwise around the tree. To construct the edges of Q , we first add an extra vertex ∂ outside the tree, and we connect each corner of the tree τ with label 1 to the vertex ∂ by an edge starting from this corner. Then every corner of τ with label $k \geq 2$ is connected by an edge to the next corner (in cyclic ordering) with label $k - 1$. The construction can be made in a unique way so that edges do not cross and do not cross the edges of the tree. The resulting collection of edges forms a quadrangulation Q whose vertex set is $V(Q) = V(\tau) \cup \{\partial\}$. This quadrangulation is rooted at the edge connecting the first corner of the root of τ to ∂ , which is oriented so that ∂ is the root vertex. See Fig. 2 for an example.

The previous construction yields a bijection from the set of all well-labeled trees with a fixed number n of edges onto the set \mathcal{M}_n^4 of all rooted quadrangulations with n faces. This bijection is called the Cori-Vauquelin-Schaeffer bijection (the CVS bijection in short). Furthermore, the following important additional property holds. If $(\tau, (\ell_v)_{v \in V(\tau)})$ is a well-labeled tree and Q is the associated quadrangulation defined as above, then, for every $v \in V(Q) \setminus \{\partial\}$,

$$d_{\text{gr}}^Q(\partial, v) = \ell_v. \quad (2)$$

In other words, distances from the root vertex in the quadrangulation Q are given

by labels on the tree associated with Q via the CVS bijection. There is no similar expression for $d_{\text{gr}}^Q(v, v')$ when v and v' are two vertices other than ∂ , but the following upper bound turns out to be important for our purposes. For every $v, v' \in V(Q) \setminus \{\partial\}$,

$$d_{\text{gr}}^Q(v, v') \leq \ell_v + \ell_{v'} - 2 \max \left(\min_{w \in [v, v']} \ell_w, \min_{w \in [v', v]} \ell_w \right) + 2, \quad (3)$$

where $[v, v']$ stands for the set of all vertices visited when going from v to v' in clockwise order around the tree (for instance, in the tree of Fig. 2, if $v = 111$ and $v' = 12$, $[v, v'] = \{111, 11, 112, 1, 12\}$ and $[v', v] = \{12, 1, \emptyset, 11, 111\}$). The bound (3) is easily derived by the following argument. Consider any corner of v . Via the CVS bijection, this corner is connected by an edge of Q to a corner of another vertex v_1 with label $\ell_v - 1$, and then this corner of v_1 is connected by an edge of Q to a corner of a second vertex v_2 with label $\ell_v - 2$. Recalling that labels correspond to distances from ∂ , we get a geodesic path $\gamma = (v, v_1, v_2, \dots)$ from v to ∂ . We may construct a similar geodesic path γ' from v' to ∂ and observe that the geodesic paths γ and γ' will eventually merge. Considering the path from v to v' obtained by concatenating the parts of γ and γ' before their merging point (and choosing the initial corners in an optimal way) easily leads to the bound (3).

The basic underlying idea of our construction of the Brownian map in the next section is to use a continuous analog of the CVS bijection. In this analog, the role of plane trees will be played by Aldous' continuum random tree, which is known to be the (universal) scaling limit of many different classes of random discrete trees. We refer to [34] for a discussion of scaling limits of labeled plane trees.

3. The Brownian map

3.1. The Gromov-Hausdorff distance. Let (E_1, d_1) and (E_2, d_2) be two compact metric spaces. The Gromov-Hausdorff distance between (E_1, d_1) and (E_2, d_2) is

$$d_{GH}(E_1, E_2) = \inf \left(d_{\text{Haus}}(\varphi_1(E_1), \varphi_2(E_2)) \right),$$

where the infimum is over all isometric embeddings $\varphi_1 : E_1 \rightarrow E$ and $\varphi_2 : E_2 \rightarrow E$ of E_1 and E_2 into the same metric space (E, d) , and d_{Haus} stands for the usual Hausdorff distance between compact subsets of E . If \mathbb{K} denotes the space of all isometry classes of compact metric spaces, then d_{GH} is a distance on \mathbb{K} , and moreover the metric space (\mathbb{K}, d_{GH}) is Polish, that is, separable and complete. We refer to Chapter 7 of Burago, Burago and Ivanov [16] for a thorough discussion of the Gromov-Hausdorff distance.

3.2. The main theorem. Recall our notation \mathcal{M}_n^p for the space of all rooted p -angulations with n faces. Note that when p is odd we consider only even values of n . With every $M \in \mathcal{M}_n^p$, we can associate the metric space $(V(M), d_{\text{gr}}^M)$.

Theorem 3.1 ([31]). *Suppose that either $p = 3$ or $p \geq 4$ is an even integer, and set*

$$c_p := \left(\frac{9}{p(p-2)} \right)^{1/4}$$

if p is even, and

$$c_3 := 6^{1/4}.$$

For every integer $n \geq 1$ (for every even integer $n \geq 2$ if $p = 3$), let M_n be uniformly distributed over \mathcal{M}_n^p . There exists a random compact metric space (\mathbf{m}_∞, D) called the Brownian map, which does not depend on p , such that

$$(V(M_n), c_p n^{-1/4} d_{\text{gr}}^{M_n}) \xrightarrow[n \rightarrow \infty]{(d)} (\mathbf{m}_\infty, D)$$

where the convergence holds in distribution in the space (\mathbb{K}, d_{GH}) .

The role of the scaling constants c_p is only to ensure that the limit does not depend on p . The convergence in distribution of the theorem may be rephrased by saying that one can construct the full sequence (M_n) in such a way that the associated metric spaces $(V(M_n), c_p n^{-1/4} d_{\text{gr}}^{M_n})$ converge to (\mathbf{m}_∞, D) for the Gromov-Hausdorff distance, outside a set of zero probability. The name Brownian map is due to Marckert and Mokkadem [37], who proved a weak form of the theorem when $p = 4$. As was already pointed in the introduction, the case $p = 4$ of the theorem has been derived independently by Miermont [41] using a different approach.




Figure 3. Simulation of a large triangulation of the sphere. Here only the graph structure is represented in three dimensions.

Remark. It is very likely that the convergence of the theorem also holds for odd values of $p \geq 5$, though additional technical difficulties appear in that case. Similarly, it is expected that a version of the convergence holds for Boltzmann distributed random planar maps, such that the probability of a given planar map M (with a fixed number n of vertices) will be proportional to

$$\prod_{f \text{ face of } M} w_{\deg(f)}$$

where $(w_k)_{k \geq 1}$ is a suitable sequence of weights. In the bipartite case where $w_k = 0$ when k is odd, a version of the convergence of the theorem holds for such random planar maps [31] under appropriate assumptions. Other recent extensions of the theorem have been mentioned in the introduction above.

In the next two subsections, we will present a precise construction of the Brownian map as a quotient space of another (well-known) random compact metric space, which is the Brownian continuum random tree or CRT (see Aldous [3, 4]).

3.3. The Brownian continuum random tree. We first recall the notion of an \mathbb{R} -tree.

Definition 3.2. A metric space (\mathcal{T}, d) is an \mathbb{R} -tree if the following two properties hold for every $a, b \in \mathcal{T}$.

- (a) There is a unique isometric map $f_{a,b}$ from $[0, d(a, b)]$ into \mathcal{T} such that $f_{a,b}(0) = a$ and $f_{a,b}(d(a, b)) = b$.
- (b) If q is a continuous injective map from $[0, 1]$ into \mathcal{T} , such that $q(0) = a$ and $q(1) = b$, we have

$$q([0, 1]) = f_{a,b}([0, d(a, b)]).$$

A rooted \mathbb{R} -tree is an \mathbb{R} -tree (\mathcal{T}, d) with a distinguished vertex ρ called the root.

We will be interested mainly in compact \mathbb{R} -trees. Informally, one should think of a compact \mathbb{R} -tree as a connected union of line segments in the plane with no loops, which is equipped with the appropriate (intrinsic) metric. For any two points a and b in the tree, there is a unique arc going from a to b in the tree, which is isometric to a line segment.

Rooted \mathbb{R} -trees can be coded by contour functions, in a way very similar to the well-known coding of plane trees by Dyck paths. Let $g : [0, 1] \rightarrow \mathbb{R}_+$ be a nonnegative continuous function such that $g(0) = g(1) = 0$ and, for every $s, t \in [0, 1]$, set

$$m_g(s, t) := \inf_{r \in [s \wedge t, s \vee t]} g(r),$$

and

$$d_g(s, t) := g(s) + g(t) - 2m_g(s, t).$$

It is easy to verify that d_g is a pseudo-metric on $[0, 1]$. As usual, we introduce the equivalence relation $s \sim_g t$ if and only if $d_g(s, t) = 0$ (or equivalently if and only if

$g(s) = g(t) = m_g(s, t)$). The function d_g induces a distance on the quotient space $\mathcal{T}_g := [0, 1] / \sim_g$, and we keep the notation d_g for this distance. We also write p_g for the canonical projection from $[0, 1]$ onto \mathcal{T}_g . Then it is not hard to verify that the metric space (\mathcal{T}_g, d_g) is a compact \mathbb{R} -tree (see e.g. [25]), which by definition is rooted at $\rho_g = p_g(0) = p_g(1)$. Furthermore the mapping $g \rightarrow \mathcal{T}_g$ is continuous with respect to the Gromov-Hausdorff distance, if the set of continuous functions g is equipped with the supremum distance.

The preceding coding induces a cyclic ordering on the tree \mathcal{T}_g , and it will be important for us to consider the corresponding intervals. By convention, if $s, t \in [0, 1]$ are such that $s > t$, we set $[s, t] := [s, 1] \cup [0, t]$. Then we note that, for every $a, b \in \mathcal{T}_g$ with $a \neq b$, there exists a smallest interval $[s, t]$ such that $p_g(s) = a$ and $p_g(t) = b$, and we define $[a, b] := p_g([s, t])$. Roughly speaking, $[a, b]$ corresponds to the set of vertices that are visited when going from a to b in “clockwise order around the tree”.

Let $\mathbf{e} = (\mathbf{e}_t)_{0 \leq t \leq 1}$ be a normalized Brownian excursion. Informally, \mathbf{e} behaves like a linear Brownian started from 0 at time 0, which is conditioned to stay positive over the time interval $(0, 1)$ and to return to 0 at time 1 (of course these conditionings require special care, see e.g. [43, Chapter XII] for a rigorous definition and many properties of the Brownian excursion).

Definition 3.3. *The CRT is the random \mathbb{R} -tree $(\mathcal{T}_{\mathbf{e}}, d_{\mathbf{e}})$ coded by the Brownian excursion \mathbf{e} .*




Figure 4. A large discrete random tree, which is an approximation of the CRT $\mathcal{T}_{\mathbf{e}}$.

3.4. Constructing the Brownian map. Analogously to the construction of quadrangulations via the CVS bijection, we will need to introduce labels on \mathbb{R} -trees. Consider first a deterministic \mathbb{R} -tree (\mathcal{T}, d) , which is rooted at ρ . We define Brownian motion indexed by \mathcal{T} as the centered Gaussian process $(\mathcal{Z}_a)_{a \in \mathcal{T}}$ whose distribution is characterised by the properties $\mathcal{Z}_\rho = 0$ and

$$E[(\mathcal{Z}_a - \mathcal{Z}_b)^2] = d(a, b)$$

for every $a, b \in \mathcal{T}$. For a general \mathbb{R} -tree, $(\mathcal{Z}_a)_{a \in \mathcal{T}}$ needs not have continuous paths. However, if $\mathcal{T} = \mathcal{T}_g$ and the function g is Hölder continuous, then it is easy to verify that we can construct $(\mathcal{Z}_a)_{a \in \mathcal{T}_g}$ so that it has continuous sample paths. This applies in particular to (almost every realization of) the CRT \mathcal{T}_e .

The building blocks of our construction are first the CRT (\mathcal{T}_e, d_e) and then the process $(Z_a)_{a \in \mathcal{T}_e}$ which, conditionally on the CRT, is distributed as Brownian motion indexed by \mathcal{T}_e . Note that there seems to be a technical difficulty here since we are considering a random process Z indexed by a random set \mathcal{T}_e : A simple way out is first to construct a process $(Z_t)_{t \in [0,1]}$ as the “tip” of the Brownian snake driven by e (see [28]), and then to observe that $Z_s = Z_t$ if $s \sim_e t$, which allows one to view Z as indexed by \mathcal{T}_e .

We set, for every $a, b \in \mathcal{T}_e$,

$$D^\circ(a, b) := Z_a + Z_b - 2 \max \left(\min_{c \in [a, b]} Z_c, \min_{c \in [b, a]} Z_c \right). \quad (4)$$

Note that we are using the notion of a tree interval which was introduced in the previous subsection, and that $D^\circ(a, b)$ is a continuous analog of the right side of (3).

The function D° is symmetric but does not satisfy the triangle inequality. We then consider the largest symmetric function that is bounded above by D° and satisfies the triangle inequality: for every $a, b \in \mathcal{T}_e$,

$$D(a, b) := \inf \left\{ \sum_{i=1}^k D^\circ(a_{i-1}, a_i) \right\}, \quad (5)$$

where the infimum is over all choices of the integer $k \geq 1$ and of the elements $a = a_0, a_1, \dots, a_{k-1}, a_k = b$ of \mathcal{T}_e .

Then D is a pseudo-metric on \mathcal{T}_e , and we set

$$a \approx b \quad \text{if and only if} \quad D(a, b) = 0.$$

Definition 3.4. *The Brownian map is the quotient space $\mathbf{m}_\infty := \mathcal{T}_e / \approx$ equipped with the distance induced by D .*

We will write Π for the canonical projection from \mathcal{T}_e onto \mathbf{m}_∞ , and we keep the notation D for the induced distance on \mathbf{m}_∞ , so that $D(\Pi(a), \Pi(b)) = D(a, b)$ for every $a, b \in \mathcal{T}_e$.

One can prove [29] that the property $D(a, b) = 0$ holds if and only if $D^\circ(a, b) = 0$, or equivalently

$$Z_a = Z_b = \max \left(\min_{c \in [a, b]} Z_c, \min_{c \in [b, a]} Z_c \right). \quad (6)$$

In other words, two vertices of the CRT are identified if they have the same label $Z_a = Z_b$ and if one can go from a to b “around the tree” (in clockwise or counter-clockwise order) encountering only vertices with label greater than or equal to Z_a . In a sense, not many pairs of vertices are identified. Only leaves of the CRT may be identified (a leaf is a vertex a of \mathcal{T}_e such that $\mathcal{T}_e \setminus \{a\}$ remains connected), and the set of all vertices a that are identified to another one has Hausdorff dimension 1 (whereas the CRT had dimension 2). Furthermore, there is only a countable collection of equivalence classes of \approx that contain 3 points, and no equivalence class contains more than 3 points. Still these identifications drastically change the topology.

Remark. As already mentioned above, the preceding construction of the Brownian map is analogous to the construction of quadrangulations via well-labeled trees in the CVS bijection (and to the construction of more general planar maps using other similar bijections). The pair $(\mathcal{T}_e, (Z_a)_{a \in \mathcal{T}_e})$ is a kind of continuous analog of a well-labeled tree $(\tau, (\ell_v)_{v \in V(\tau)})$ (there is a minor difference because we do not impose a positivity condition on the labels in the continuous setting, but we could have re-rooted \mathcal{T}_e at the vertex with minimal label and shifted all labels Z_a so that they become nonnegative, which would not have affected the preceding construction). Still one may notice that no identification of vertices is needed in the discrete setting. The reason why such identifications become necessary in the continuous setting can be explained intuitively as follows. In a large well-labeled tree, there will exist vertices u and v , which are at a “macroscopic distance” in the tree, and such that $\ell_v = \ell_u - 1$ and $\ell_w \geq \ell_u$ for every vertex $w \in [u, v]$, where the interval $[u, v]$ is as in (3) (note that (6) is just a continuous version of these properties): According to the rules of the CVS bijection, two such vertices are linked by an edge of the quadrangulation, and asymptotically, recalling that the graph distance is rescaled by a factor tending to 0, this leads to an identification of two vertices in the tree.

3.5. Properties of the Brownian map. In this subsection, we describe a few properties of the Brownian map, which show that this random metric space behaves very differently from a smooth surface.

Proposition 3.5. (i) *The Brownian map (\mathbf{m}_∞, D) is a.s. homeomorphic to the sphere \mathbb{S}^2 .*

(ii) *The Hausdorff dimension of (\mathbf{m}_∞, D) is a.s. equal to 4.*

Part (i) is proved in [35], and part (ii) was derived in [29]. In view of Proposition 3.5, one may look at the Brownian map as a sphere with a very singular metric. Property (ii) is of course reminiscent of the fact that the Hausdorff dimension of a Brownian path in \mathbb{R}^d , $d \geq 2$ is twice the dimension of a smooth curve.

We next turn to uniform estimates on the volume of balls. We first need to introduce the volume measure on \mathbf{m}_∞ , which in a sense is uniformly distributed over this random space. We set $\mathbf{p} := \Pi \circ p_e$, which maps $[0, 1]$ onto \mathbf{m}_∞ . The volume measure Λ on the Brownian map is the image of Lebesgue measure on $[0, 1]$ under \mathbf{p} . We could also have introduced this volume measure via the convergence

in Theorem 3.1, as the limit of the normalized counting measure on the vertex set of M_n (to make this rigorous, one should state the latter convergence in the sense of the Gromov-Hausdorff topology for measured metric spaces).

If $x \in \mathbf{m}_\infty$ and $r > 0$, we write $B_D(x, r)$ for the closed ball of radius r centered at x in (\mathbf{m}_∞, D) .

Proposition 3.6. *Let $\delta > 0$. There exist two (random) positive constants c_δ and C_δ such that, for every $r \in (0, 1]$ and every $x \in \mathbf{m}_\infty$,*

$$c_\delta r^{4+\delta} \leq \Lambda(B_D(x, r)) \leq C_\delta r^{4-\delta}.$$

This shows that in a way the Brownian map is “very regular in its irregularity”: The volume of any small ball of radius r is approximately r^4 . The upper bound in Proposition 3.6 is proved in [30], and the lower bound is easy from the bound $D \leq D^\circ$ and the Hölder continuity properties of the process Z .

Another way of quantifying the irregularity of the Brownian map is to look at connected components of the complement of a ball. If this complement is not empty, it will have typically infinitely many connected components. The following proposition from [32] gives a more precise estimate.

Proposition 3.7. *Suppose that U is a point of \mathbf{m}_∞ chosen at random according to the volume measure Λ . Let $r > 0$ and, for every $\varepsilon > 0$, let N_ε^r be the number of connected components of $B_D(U, r)^c$ that intersect $B_D(U, r + \varepsilon)^c$. Then,*

$$\varepsilon^3 N_\varepsilon^r \xrightarrow{\varepsilon \rightarrow 0} c_0 \lambda_r,$$

where the convergence holds in probability, $c_0 > 0$ is an explicit constant and λ_r is the value at r of the continuous density of the profile of distances from U , which is the random measure \mathcal{I} on \mathbb{R}_+ defined by

$$\mathcal{I}(A) = \int \Lambda(dx) \mathbf{1}_A(D(U, x)),$$

for any Borel subset A of \mathbb{R}_+ .

4. Geodesics in the Brownian map

Although the Brownian map remains a mysterious object in many respects, there is detailed information about the structure of geodesics toward a typical point. In this section we describe these results following [30]. We rely on the construction given in subsection 3.4.

One can prove (see in particular [36]) that there exists a unique vertex $\rho_* \in \mathcal{T}_e$ such that

$$Z_{\rho_*} = \min_{c \in \mathcal{T}_e} Z_c.$$

We set $x_* := \Pi(\rho_*)$. In what follows we discuss geodesics toward x_* in the Brownian map. It should be emphasized that x_* is not a special point of the Brownian

map: If U is a point of the Brownian map chosen at random according to the volume measure, one can check [30] that the random pointed compact metric spaces $(\mathbf{m}_\infty, D, U)$ and $(\mathbf{m}_\infty, D, x_*)$ have the same distribution, so that our results also give information about geodesics toward the “typical” point U .

To simplify notation, we set

$$Z_* := \min_{c \in \mathcal{T}_e} Z_c.$$

We first observe that, for every $a \in \mathcal{T}_e$,

$$D(x_*, \Pi(a)) = Z_a - Z_*. \quad (7)$$

Indeed, it immediately follows from (4) that $D(x_*, \Pi(a)) \leq D^\circ(\rho_*, a) = Z_a - Z_*$. On the other hand, by using the easy bound

$$D^\circ(a, b) \geq |Z_a - Z_b|, \quad (8)$$

which also follows from (4), we immediately get from (5) that the reverse inequality $D(x_*, \Pi(a)) = D(\rho_*, a) \geq Z_a - Z_*$ holds.

Let $x = \Pi(a)$, $a \in \mathcal{T}_e$ be a point in the Brownian map. We can construct certain geodesics from x to x_* in the following manner. Choose $t \in [0, 1)$ such that $p_e(t) = a$, and for every $r \in [0, Z_a - Z_*]$, set

$$\gamma_t(r) := \begin{cases} \min\{s \in [t, 1] : Z_{p_e(s)} = Z_a - r\} & \text{if } \{s \in [t, 1] : Z_{p_e(s)} = Z_a - r\} \neq \emptyset, \\ \min\{s \in [0, t] : Z_{p_e(s)} = Z_a - r\} & \text{otherwise.} \end{cases}$$

Informally, $p_e(\gamma_t(r))$ is the first vertex with label $Z_a - r$ that one encounters when exploring the tree clockwise starting from a . Note that, when a is not a leaf, there are more than one way of starting from a : This corresponds to the different possible choices of t . Since ρ_* is the unique vertex with minimal label, it is clear that $p_e(\gamma_t(Z_a - Z_*)) = \rho_*$.

Lemma 4.1. *Let $a \in \mathcal{T}_e$ and $t \in [0, 1)$ such that $p_e(t) = a$. Set $\Gamma_t(r) = \mathbf{p}(\gamma_t(r))$ for every $r \in [0, Z_a - Z_*]$. Then $(\Gamma_t(r))_{0 \leq r \leq Z_a - Z_*}$ is a geodesic from $\Pi(a)$ to x_* in (\mathbf{m}_∞, D) . Such geodesics are called simple geodesics.*

The proof is easy. If $0 \leq r \leq r' \leq Z_a - Z_*$, the bound $D \leq D^\circ$ and the definition (4) immediately shows that $D(\Gamma_t(r), \Gamma_t(r')) \leq D^\circ(p_e(\gamma_t(r)), p_e(\gamma_t(r'))) = r' - r$. On the other hand, by (7), $D(\Gamma_t(0), \Gamma_t(Z_a - Z_*)) = D(\Pi(a), x_*) = Z_a - Z_*$, and the triangle inequality now gives $D(\Gamma_t(r), \Gamma_t(r')) = r' - r$ for every $0 \leq r \leq r' \leq Z_a - Z_*$.

Simple geodesics are indeed analogs in our continuous setting of the discrete geodesics for quadrangulations that were briefly discussed at the end of Section 2. The following proposition [30] is the key result in our study of geodesics.

Proposition 4.2. *All geodesics to x_* in (\mathbf{m}_∞, D) are simple geodesics.*

This proposition makes it possible to classify all the geodesics to x_* . Indeed, it is easy to count simple geodesics. If we fix a point $x \in \mathbf{m}_\infty$, a simple geodesic γ_t from x to x_* is obtained by choosing first $a \in \mathcal{T}_e$ such that $\Pi(a) = x$, and then $t \in [0, 1)$ such that $p_e(t) = a$. The choice of a is in fact irrelevant: If $\Pi^{-1}(x)$ is not a singleton, then the vertices a in $\Pi^{-1}(x)$ must be leaves, then, for each such a , there is a unique $t \in [0, 1)$ with $p_e(t) = a$, and one immediately verifies that the associated simple geodesics coincide. On the other hand, if $x = \Pi(a)$ and a is not a leaf of \mathcal{T}_e , then there are (2 or 3) values of t such that $p_e(t) = a$, and the corresponding simple geodesics are distinct. The preceding discussion leads to the following theorem.

Theorem 4.3. *Let $\text{Sk}(\mathcal{T}_e)$ stand for the set of all vertices of the CRT that are not leaves, and $\text{Skel} := \Pi(\text{Sk}(\mathcal{T}_e))$. Then the restriction of Π to $\text{Sk}(\mathcal{T}_e)$ is a homeomorphism, and the Hausdorff dimension of Skel is equal to 2. Furthermore, a.s. for every $x \in \mathbf{m}_\infty$,*

- *if $x \in \mathbf{m}_\infty \setminus \text{Skel}$, there is a unique geodesic from x to x_* ;*
- *if $x \in \text{Skel}$, the number of distinct geodesics from x to x_* is the multiplicity of x in Skel , that is, the number of connected components of $\text{Skel} \setminus \{x\}$. This multiplicity is either 2 or 3.*

The set Skel , which is a dense subset of the Brownian map homeomorphic to a non-compact real tree, thus appears as the cut-locus of the Brownian map with respect to the point x_* : A point x belongs to Skel if and only if there are at least two distinct geodesics from x to x_* . Perhaps suprisingly, there is a strong analogy with classical results of differential geometry that go back to Poincaré. For a smooth surface homeomorphic to the sphere, the cut-locus is also a tree, and the number of distinct geodesics from a point of the cut-locus is equal to its multiplicity.

It is easy to verify that the set Skel has zero volume measure (in terms of Hausdorff dimension, it is already clear that Skel is a “small” subset of \mathbf{m}_∞). A consequence of Theorem 4.3 is thus the fact that, if one picks independently two points x and x' according to the volume measure on \mathbf{m}_∞ , then a.s. there is a unique geodesic between x and x' (see also Miermont [40] for a related result with a different approach). From this and Theorem 3.1, one can deduce a property of “macroscopic uniqueness” of discrete geodesics in large planar maps, which was already mentioned in the introduction. See [30] for more details.

Theorem 4.3 shows that our construction of the Brownian map as a quotient space of the CRT has a strong geometric meaning (although it is certainly not the only possible construction). Indeed the set Skel , which is a homeomorphic image of the skeleton of the CRT in our construction, is a geometric object defined intrinsically in terms of the Brownian map, and thus does not depend on the particular construction we have developed.

Remark. The version of Theorem 4.3 given in [30] applies to any random compact metric space which appears as a Gromov-Hausdorff limit in distribution of rescaled

$2p$ -angulations (the uniqueness of such a limit was not yet known). This description of geodesics was then a key ingredient of the proof of Theorem 3.1 in [31]

The following corollary gives a confluence property of geodesics, which easily follows from the fact that two simple geodesics will always merge before their endpoint.

Corollary 4.4. *Let $\delta > 0$. Then a.s. there exists $\varepsilon > 0$ such that, whenever x and x' are two points of \mathbf{m}_∞ with $D(x_*, x) \geq \delta$ and $D(x_*, x') \geq \delta$, then, if f is a geodesic from x_* to x and g is a geodesic from x_* to y , we have $f(r) = g(r)$ for every $r \in [0, \varepsilon]$.*

Informally, there is only one way of leaving x_* along a geodesic. By the remarks of the beginning of this section, the same property holds for a typical point chosen according to the volume measure on \mathbf{m}_∞ .

5. Infinite volume limits and the Brownian plane

5.1. UIPT and UIPQ. Theorem 3.1 deals with the convergence of uniformly distributed rooted p -angulations with n faces in the Gromov-Hausdorff sense, provided that the graph distance is rescaled by the factor $n^{-1/4}$ when n tends to infinity. Note that this rescaling is necessary if we want to obtain a compact limit. On the other hand, one may also consider the convergence of the same random objects without rescaling, but then in a different sense than the Gromov-Hausdorff convergence. This leads to infinite random lattices.

We consider possibly infinite (multi)graphs that are always connected, pointed (meaning that there is a distinguished vertex ρ called the root vertex) and locally finite in the sense that the degree of every vertex is finite. As previously, these graphs are equipped with the graph distance. A ball of radius k centered at a vertex v of the graph is then viewed as the subgraph consisting of all vertices at distance less than or equal to k from v and the edges connecting these vertices. A sequence (G_n, ρ_n) of pointed graphs is said to converge locally to a limiting pointed graph (G, ρ) if for every integer $k \geq 0$, for every n sufficiently large, the ball of radius k centered at ρ_n in G_n is equal to the ball of radius k centered at ρ in G (equality here is in the sense of isomorphism between finite graphs with a distinguished vertex).

The following result is due to Angel and Schramm [7] in the case of triangulations and to Krikun [27] (see also [23]) in the case of quadrangulations.

Theorem 5.1. *Let $p = 3$ or $p = 4$, and let M_n be uniformly distributed over \mathcal{M}_n^p . There exists a random infinite graph $M_\infty^{(p)}$ such that*

$$M_n \xrightarrow[n \rightarrow \infty]{(d)} M_\infty^{(p)},$$

where the convergence holds in distribution in the sense of the local convergence of graphs.

In this theorem, the convergence just means that, for every integer $k \geq 0$, the probability that the ball of radius k centered at the root vertex in M_n is equal to a given graph converges to the same probability for the limit $M_\infty^{(p)}$. A completely different construction of $M_\infty^{(4)}$ based on a version of the CVS bijection for infinite trees was given in [17] (see [38] for a proof of the fact that the two constructions give the same object).

Remark. One may also define the limit in Theorem 5.1 as an infinite planar map, that is, with a given embedding in the sphere, and this makes it possible to give a stronger form of the local convergence (see [7] and [27]). Here for simplicity we avoid dealing with infinite planar maps, since only the graph structure will play a role in the subsequent statements.

The infinite random graph $M_\infty^{(p)}$ is called the uniform infinite planar triangulation (UIPT) when $p = 3$ and the uniform infinite planar quadrangulation (UIPQ) when $p = 4$. Properties of these infinite random graphs have been investigated in detail in the recent years. In particular, Angel [6] has studied percolation on the UIPT. The recurrence of simple random walk on the UIPT or the UIPQ has been obtained recently by Gurel-Gurevich and Nachmias [26]. See also Benjamini and Curien [10] for a proof of subdiffusivity of simple random walk on the UIPQ.

5.2. Convergence to the Brownian plane. The following results are taken from [21]. To simplify notation, we write $Q_\infty = M_\infty^{(4)}$ for the UIPQ, which was introduced in the previous subsection. The next theorem shows that, if we rescale the graph distance on the UIPQ by a factor tending to 0, the resulting metric spaces converge (in a suitable sense) toward a limiting random non-compact metric space, which is called the Brownian plane. We recall that a pointed metric space is just a metric space equipped with a distinguished point.

Theorem 5.2. *Let $V(Q_\infty)$ denote the vertex set of Q_∞ , which is equipped with the graph distance d_{gr} , and let ρ_{Q_∞} stand for the root vertex of Q_∞ . There exists a random non-compact pointed metric space $(\mathcal{P}, D_\infty, \rho_\infty)$ called the Brownian plane such that*

$$(V(Q_\infty), \lambda d_{\text{gr}}, \rho_{Q_\infty}) \xrightarrow[\lambda \rightarrow 0]{(d)} (\mathcal{P}, D_\infty, \rho_\infty),$$

where the convergence holds in distribution in the local Gromov-Hausdorff sense.

We refer to [16] for a precise definition of the local Gromov-Hausdorff convergence for (non-compact) pointed metric spaces. In the present setting, this means that, for every $r > 0$, the ball of radius r centered at ρ_{Q_∞} in $(V(Q_\infty), \lambda d_{\text{gr}})$ will converge in distribution in the Gromov-Hausdorff sense to the ball of radius r centered at ρ_∞ in the limiting space (\mathcal{P}, D_∞) – in fact we should use here the Gromov-Hausdorff distance between pointed compact metric spaces, which is defined by a minor modification of the definition in subsection 3.1.

The Brownian plane $(\mathcal{P}, D_\infty, \rho_\infty)$ appears in several other limit theorems. In particular, if in the convergence of Theorem 3.1 for $p = 4$ one replaces the scaling factor $n^{-1/4}$ by a function $\beta(n)$ tending to 0 but such that $n^{1/4}\beta(n)$ tends to

infinity, the convergence still holds in the local Gromov-Hausdorff sense and the limit is now the Brownian plane.

Alternatively, the Brownian plane can be viewed as the tangent cone in distribution of the Brownian map at a distinguished vertex U chosen at random according to the volume measure Λ . In fact, a stronger property holds: One can construct, on the same probability space, both the Brownian map \mathbf{m}_∞ and the Brownian plane \mathcal{P} , in such a way that, a.s., there exists a (random) $\varepsilon > 0$ such that the balls of radius ε centered at the distinguished point in \mathbf{m}_∞ and in \mathcal{P} are isometric. The latter fact, together with the scale invariance property of \mathcal{P} (see subsection 5.3 below), allows one to derive many properties of the Brownian plane from those known for the Brownian map.

We will now give a precise construction of the Brownian plane. Not surprisingly, this construction is very similar to the one developed above for the Brownian map. We consider two independent three-dimensional Bessel processes R and R' started from 0 (see e.g. [43] for basic facts about Bessel processes). We then define a process $Y = (Y_t)_{t \in \mathbb{R}}$ indexed by the real line, by setting

$$Y_t := \begin{cases} R_t & \text{if } t \geq 0, \\ R'_{-t} & \text{if } t \leq 0. \end{cases}$$

Then, for every $s, t \in \mathbb{R}$, we set

$$m_Y(s, t) := \inf_{r \in \overline{st}} Y_r,$$

with the notation $\overline{st} = [s \wedge t, s \vee t]$ if $st \geq 0$, $\overline{st} = (-\infty, s \wedge t] \cup [s \vee t, \infty)$ if $st < 0$. We define a random pseudo-distance on \mathbb{R} by

$$d_Y(s, t) := Y_s + Y_t - 2m_Y(s, t)$$

and set $s \sim_Y t$ if $d_Y(s, t) = 0$. The quotient space $\mathcal{T}_\infty = \mathbb{R} / \sim_Y$ equipped with d_Y is a (non-compact) random real tree, which is sometimes called the infinite Brownian tree. We write $p_\infty : \mathbb{R} \rightarrow \mathcal{T}_\infty$ for the canonical projection and set $\rho_\infty := p_\infty(0)$, which plays the role of the root of \mathcal{T}_∞ . The volume measure on \mathcal{T}_∞ is the image of Lebesgue measure on \mathbb{R} under p_∞ .

We next consider Brownian motion indexed by \mathcal{T}_∞ . Formally, we consider a real-valued process $(Z_t^\infty)_{t \in \mathbb{R}}$ such that, conditionally given the process Y , Z^∞ is a centered Gaussian process with covariance

$$E[Z_s^\infty Z_t^\infty | Y] = m_Y(s, t),$$

so that we have $Z_0^\infty = 0$ and $E[(Z_s^\infty - Z_t^\infty)^2 | Y] = d_Y(s, t)$. It is not hard to verify that the process Z^∞ has a modification with continuous paths. Then a.s. we have $Z_s^\infty = Z_t^\infty$ for every $s, t \in \mathbb{R}$ such that $d_Y(s, t) = 0$ and therefore we may view Z^∞ as indexed by \mathcal{T}_∞ .

For every $s, t \in \mathbb{R}$, we set

$$D_\infty^\circ(s, t) := Z_s^\infty + Z_t^\infty - 2 \min_{r \in [s \wedge t, s \vee t]} Z_r^\infty.$$

We extend the definition of D_∞° to $\mathcal{T}_\infty \times \mathcal{T}_\infty$ by setting for $a, b \in \mathcal{T}_\infty$,

$$D_\infty^\circ(a, b) := \min\{D_\infty^\circ(s, t) : s, t \in \mathbb{R}, p_\infty(s) = a, p_\infty(t) = b\}.$$

Finally, we set, for every $a, b \in \mathcal{T}_\infty$,

$$D_\infty(a, b) := \inf_{a_0=a, a_1, \dots, a_k=b} \sum_{i=1}^k D_\infty^\circ(a_{i-1}, a_i)$$

where the infimum is over all choices of the integer $k \geq 1$ and of the finite sequence a_0, a_1, \dots, a_k in \mathcal{T}_∞ such that $a_0 = a$ and $a_k = b$. Then D_∞ is a pseudo-distance on \mathcal{T}_∞ , and we set $a \approx b$ if $D_\infty(a, b) = 0$ (this can be proved to be equivalent to the property $D_\infty^\circ(a, b) = 0$). The Brownian plane is the quotient space $\mathcal{P} = \mathcal{T}_\infty / \approx$, which is equipped with the metric induced by D_∞ and with the distinguished point which is the equivalence class of ρ_∞ (without risk of confusion, we still write ρ_∞ for this equivalence class). The volume measure on \mathcal{P} is the image of the volume measure on \mathcal{T}_∞ under the canonical projection.

5.3. Properties of the Brownian plane. The Brownian plane is scale invariant, meaning that, for every $\lambda > 0$, the space $(\mathcal{P}, \lambda D_\infty, \rho_\infty)$ has the same distribution as $(\mathcal{P}, D_\infty, \rho_\infty)$. This property can be derived from the explicit construction given above, or from the fact that the Brownian plane is a tangent cone in distribution to the Brownian map.

Other properties of the Brownian plane are very similar to those of the Brownian map. We state an analog of Proposition 3.5.

Proposition 5.3. *Almost surely, the Brownian plane (\mathcal{P}, D_∞) is homeomorphic to the Euclidean plane \mathbb{R}^2 and has Hausdorff dimension 4.*

From the construction of the previous subsection, the points of the Brownian plane inherit labels Z_x^∞ from the corresponding labels on the infinite Brownian tree \mathcal{T}_∞ . These labels can be interpreted as “distances from infinity” in the following sense. For every $x, y \in \mathcal{P}$,

$$Z_x^\infty - Z_y^\infty = \lim_{z \rightarrow \infty} (D_\infty(x, z) - D_\infty(y, z)).$$

The existence of the preceding limit is related to a property of confluence of geodesic rays in the Brownian plane. Recall that a geodesic ray is a continuous path $\gamma : [0, \infty) \rightarrow \mathcal{P}$ such that $D_\infty(\gamma(t), \gamma(t')) = |t - t'|$ for every $t, t' \geq 0$.

Proposition 5.4. *Let γ and γ' be two geodesic rays in \mathcal{P} . Then there exists two reals α and $\beta \geq |\alpha|$ such that $\gamma(t) = \gamma'(\alpha + t)$ for every $t \geq \beta$.*

This is a kind of version at infinity of Corollary 4.4. Note that an exact analog of Corollary 4.4 also holds for the Brownian plane, and is easy to prove by the coupling argument mentioned above.

The Brownian plane seems to be more tractable than the Brownian map for explicit calculations, partly because of its scale invariance. Let us discuss some

recent results from [22] that shed light on the probabilistic structure of the Brownian plane. For every $r > 0$, we let B_r be the closed ball of radius r centered at ρ_∞ in \mathcal{P} , and we define the “extended ball” B_r^\bullet as the complement of the unbounded connected component of $(B_r)^c$. Informally, B_r^\bullet is obtained by “filling the holes” in B_r . We write $|B_r^\bullet|$ for the volume of B_r^\bullet .

Proposition 5.5. *For every $\lambda > 0$,*

$$E[\exp(-\lambda|B_r^\bullet|)] = \frac{3^{3/2} \cosh((2\lambda)^{1/4}r)}{\left(\cosh^2((2\lambda)^{1/4}r) + 2\right)^{3/2}}.$$

In fact, one can give a simple description of the whole process $(|B_r^\bullet|)_{r \geq 0}$. Set $\psi(\lambda) := (\frac{8}{3})^{1/2} \lambda^{3/2}$ and recall that the continuous-state branching process with branching mechanism ψ is the Markov process in \mathbb{R}_+ whose transition kernels are characterized by the Laplace transform

$$E[\exp(-\lambda X_t) \mid X_0 = x] = \exp(-x u_t(\lambda))$$

where the function $u_t(\lambda)$ solves the differential equation $\frac{du_t(\lambda)}{dt} = -\psi(u_t(\lambda))$ with initial condition $u_0(\lambda) = \lambda$ (see e.g. [28]). Note that the sample paths of X are right-continuous with left limits and that X is absorbed at 0. We can define a process $X^\infty = (X_t^\infty)_{t \leq 0}$, which is indexed by negative times and corresponds to X “started from $+\infty$ at time $-\infty$ and conditioned to hit 0 at time 0” (formally, we start X at time 0 from $x > 0$, we then shift time so that the hitting time of 0 becomes 0, and we finally let x tend to $+\infty$).

Proposition 5.6. *The process $(|B_r^\bullet|)_{r \geq 0}$ has the same distribution as the process $(W_r)_{r \geq 0}$ defined by*

$$W_r = \sum_{-r \leq u \leq 0} \xi_u (\Delta X_u^\infty)^2,$$

where ΔX_u^∞ stands for the jump of X^∞ at time u , and, conditionally given X^∞ , the nonnegative random variables ξ_u are independent and identically distributed with density

$$\frac{1}{\sqrt{2\pi}} x^{-5/2} e^{-1/2x}.$$

Informally, each jump of the process $(|B_r^\bullet|)_{r \geq 0}$ corresponds to the creation of a new connected component of $(B_r)^c$ (there are many such components, as shown by Proposition 3.7 in the Brownian map case), noting that this newly created connected component will be “swallowed” by the extended ball. For $r \geq 0$, the random variable X_r^∞ should be interpreted as the length, in a generalized sense, of the boundary of B_r^\bullet . Note that, when a connected component is swallowed, the length of the boundary of B_r^\bullet has a negative jump. The distribution of the variables ξ_u then corresponds to the law of the volume of a connected component given that its boundary has length 1.

The preceding interpretation is closely related to some results of Krikun [27] in the discrete setting of the UIPQ. The description of the process $(|B_r^\bullet|)_{r \geq 0}$ can also be interpreted in terms of the “peeling process” of the UIPT studied in [6].

6. Canonical embeddings and open questions

In this section, we come back to the questions that were discussed at the beginning of the introduction above. We would like to have a “canonical” construction of a random metric Δ on the sphere \mathbb{S}^2 , in such a way that

$$(\mathbf{m}_\infty, D) \stackrel{(d)}{=} (\mathbb{S}^2, \Delta).$$

Furthermore we expect Δ to behave well under the conformal transformations of the sphere.

Recall that the embedding of a planar map is defined up to orientation-preserving homeomorphisms of the sphere. However, there are (almost) canonical ways of choosing these embeddings. Consider the case of simple triangulations, that is, triangulations without loops or multiple edges. According the circle packing theorem, any such triangulation can be represented via a circle packing of the sphere, in such a way that the vertex set of the triangulation is the set of centers of all circles, and two vertices are linked by an edge if and only if the associated circles are tangent (see Fig. 5 for an example). This representation is in fact unique up to the conformal transformations of the sphere (the Möbius transformations).

Next suppose that, for every even integer $n \geq 2$, we have constructed a circle-packing embedding \mathcal{C}_n of a uniformly distributed simple triangulation with n faces. Write $V(\mathcal{C}_n)$ for the vertex set of \mathcal{C}_n and d_{gr}^n for the graph distance on $V(\mathcal{C}_n)$. By Theorem 3.1, or more precisely by the extension of Theorem 3.1 to simple triangulations found in [2], we have

$$\left(V(\mathcal{C}_n), \left(\frac{3}{2}\right)^{1/4} n^{-1/4} d_{\text{gr}}^n \right) \xrightarrow[n \rightarrow \infty]{(d)} (\mathbf{m}_\infty, D)$$

in the Gromov-Hausdorff sense. Note that the constant $(3/2)^{1/4}$ is different from the constant $6^{1/4}$ in (1) or in Theorem 3.1, because we are dealing with simple triangulations.

Conjecture. *One can construct the circle packing embeddings \mathcal{C}_n in such a way that*

$$\sup_{x \in \mathbb{S}^2} \left(\min_{y \in V(\mathcal{C}_n)} |x - y| \right) \xrightarrow[n \rightarrow \infty]{} 0$$

in probability, and there exists a continuous random process $(\Delta(x, y))_{x, y \in \mathbb{S}^2}$, which is nonzero outside the diagonal and such that

$$\sup_{x, y \in V(\mathcal{C}_n)} \left| \Delta(x, y) - \left(\frac{3}{2}\right)^{1/4} n^{-1/4} d_{\text{gr}}^n(x, y) \right| \xrightarrow[n \rightarrow \infty]{} 0$$

in probability.

If the conjecture holds, then it is not hard to verify that Δ defines a random metric on the sphere and that (\mathbb{S}^2, Δ) gives a representation of the Brownian map. We note that there are other ways of defining canonical embeddings of planar maps, which would lead to other versions of the conjecture. In particular, we can




Figure 5. A circle packing of the sphere associated with a simple triangulation.

associate with a triangulation of the sphere a Riemann surface, which is obtained by viewing each face as an equilateral triangle with sides of length 1 and then gluing two adjacent triangles along their common edge. The resulting Riemann surface is homeomorphic to the sphere and thus the uniformization theorem yields a canonical embedding, again modulo Möbius transformations. See [20] for an application of these ideas to the conformal structure of random planar maps.

As a final remark, we note that a very recent work of Miller and Sheffield [42] has introduced a new growth process called Quantum Loewner Evolution (QLE), which might provide a direct construction of the Brownian map, or perhaps of the Brownian plane discussed in the previous section, with conformal invariance properties. Assuming that this construction goes through, one may expect new fascinating connections between the Brownian map on one hand, the Schramm-Loewner evolutions and the two-dimensional Gaussian free field on the other hand.

References

- [1] Abraham, C. Rescaled bipartite planar maps converge to the Brownian map. Preprint (2013), available at: [arXiv:1312.5959](https://arxiv.org/abs/1312.5959)
- [2] Addario-Berry, L., Albenque, M. The scaling limit of random simple triangulations and random simple quadrangulations. Preprint (2013), available at: [arXiv:1306.5227](https://arxiv.org/abs/1306.5227)
- [3] Aldous, D., The continuum random tree I. *Ann. Probab.* **19** (1991), 1-28.

- [4] Aldous, D., The continuum random tree III. *Ann. Probab.* **21** (1993), 248-289.
- [5] Ambjørn, J., Durhuus, B., Jonsson, T., *Quantum Geometry. A Statistical Field Theory Approach*. Cambridge Monographs on Mathematical Physics. Cambridge Univ. Press, Cambridge, 1997.
- [6] Angel, O., Growth and percolation on the uniform infinite planar triangulation. *Geomet. Funct. Anal.* **3** (2003), 935-974.
- [7] Angel, O., Schramm, O. Uniform infinite planar triangulations. *Comm. Math. Phys.* **241** (2003), 191-213.
- [8] Beltran, J., Le Gall, J.-F. Quadrangulations with no pendant vertices. *Bernoulli* **19** (2013), 1150-1175.
- [9] Benjamini, I. Random planar metrics. Proceedings of the International Congress of Mathematicians. Vol. IV, 2177-2187, Hindustan Book Agency, New Delhi, 2010.
- [10] Benjamini, I., Curien, N. Simple random walk on the uniform infinite planar quadrangulation: subdiffusivity via pioneer points. *Geom. Funct. Anal.* **23** (2013), 501-531.
- [11] Bettinelli, J., Jacob, E., Miermont, G. The Scaling Limit of Uniform Random Plane Maps, via the Ambjørn-Budd Bijection. Preprint (2013), available at: arXiv:1312.5842
- [12] Bouttier, J., Di Francesco, P., Guitter, E., Planar maps as labeled mobiles. *Electronic J. Combinatorics* **11** (2004), #R69.
- [13] Bouttier, J., Guitter, E., Statistics in geodesics in large quadrangulations. *J. Phys. A* **41** (2008), 145001, 30 pp.
- [14] Bouttier, J., Guitter, E., The three-point function of planar quadrangulations. *J. Stat. Mech. Theory Exp.* (2008), P07020.
- [15] Bouttier, J., Guitter, E., Confluence of geodesic paths and separating loops in large planar quadrangulations. *J. Stat. Mech. Theory Exp.* (2009), P03001.
- [16] Burago, D., Burago, Y., Ivanov, S., *A Course in Metric Geometry*. Graduate Studies in Mathematics, vol. 33. Amer. Math. Soc., Boston, 2001.
- [17] Chassaing, P., Durhuus, B. Local limit of labeled trees and expected volume growth in a random quadrangulation. *Ann. Probab.* **34** (2006), 879-917.
- [18] Chassaing, P., Schaeffer, G., Random planar lattices and integrated superBrownian excursion. *Probab. Th. Rel. Fields* **128** (2004), 161-212.
- [19] Cori, R., Vauquelin, B., Planar maps are well labeled trees. *Canad. J. Math.*, **33** (1981), 1023-1042.
- [20] Curien, N. A glimpse of the conformal structure of random planar maps. Preprint (2013), available at: arXiv:1308.1807
- [21] Curien, N., Le Gall, J.-F., The Brownian plane. *J. Theoret. Probab.*, to appear.
- [22] Curien, N., Le Gall, J.-F., in preparation.
- [23] Curien, N., Ménard, L., Miermont, G., A view from infinity of the uniform infinite quadrangulation. *ALEA Lat. Am. J. Probab. Math. Stat.* **10** (2013), 4588.
- [24] Duplantier, B., Sheffield, S., Liouville quantum gravity and KPZ. *Invent. Math.*, **185** (2011), 333-393.
- [25] Duquesne, T., Le Gall, J.-F., Probabilistic and fractal aspects of Lévy trees. *Probab. Th. Rel. Fields* **131** (2005), 553-603.

- [26] Gurel-Gurevich, O., Nachmias, A., Recurrence of planar graph limits. *Ann. Math.* **177** (2013), 761781.
- [27] Krikun, M., Local structure of random quadrangulations. Preprint, math:PR/0512304
- [28] Le Gall, J.-F. *Spatial branching processes, random snakes and partial differential equations*. Lectures in Mathematics ETH Zürich. Birkhäuser, Basel, 1999.
- [29] Le Gall, J.-F., The topological structure of scaling limits of large planar maps. *Invent. Math.* **169** (2007), 621-670.
- [30] Le Gall, J.-F., Geodesics in large planar maps and in the Brownian map. *Acta Math.* **205** (2010), 287-360.
- [31] Le Gall, J.-F., Uniqueness and universality of the Brownian map. *Ann. Probab.* **41** (2013), 2880-2960.
- [32] Le Gall, J.-F., The Brownian cactus II. Upcrossings and local times of super-Brownian motion. Preprint, arXiv:1308.6762
- [33] Le Gall, J.-F., Miermont, G., Scaling limits of random planar maps with large faces. *Ann. Probab.* **39** (2011), 1-69.
- [34] Le Gall, J.-F., Miermont, G., Scaling limits of random trees and planar maps. Probability and statistical physics in two and more dimensions, pp. 155211, Clay Math. Proc., 15. Amer. Math. Soc., Providence, 2012.
- [35] Le Gall, J.-F., Paulin, F., Scaling limits of bipartite planar maps are homeomorphic to the 2-sphere. *Geomet. Funct. Anal.* **18** (2008), 893-918.
- [36] Le Gall, J.-F., Weill, M. Conditioned Brownian trees. *Annales Inst. H. Poincaré Probab. Stat.* **42** (2006), 455-489.
- [37] Marckert, J.-F., Mokkadem, A., Limit of normalized quadrangulations: the Brownian map. *Ann. Probab.* **34**(2006), 2144-2202.
- [38] Ménard, L., The two uniform infinite quadrangulations of the plane have the same law. *Ann. Inst. H. Poincaré Probab. Stat.* **46** (2010), 190-208.
- [39] Miermont, G., On the sphericity of scaling limits of random planar quadrangulations. *Electron. Commun. Probab.*, **13** (2008), 248-257.
- [40] Miermont, G., Tessellations of random maps of arbitrary genus. *Ann. Sci. École Norm. Sup.* **42** (2009), 725-781.
- [41] Miermont, G., The Brownian map is the scaling limit of uniform random plane quadrangulations. *Acta Math.* **210** (2013), 319-401.
- [42] Miller, J., Sheffield, S. Quantum Loewner Evolution. Preprint, arXiv:1312.5745
- [43] Revuz, D., Yor, M., *Continuous Martingales and Brownian Motion*. Springer, Berlin, 1991.
- [44] Schramm, O., Conformally invariant scaling limits: an overview and a collection of problems. Proceedings of the International Congress of Mathematicians (Madrid 2006), Vol.I, 513-543. European Math. Soc., Zürich, 2007.
- [45] Sheffield, S. Conformal weldings of random surfaces: SLE and the quantum gravity zipper. Preprint, arXiv:1012.4797
- [46] Tutte, W.T., A census of planar maps. *Canad. J. Math.*, **15** (1963), 249-271.

Université Paris-Sud, Mathématiques, Bât.425, 91405 Orsay Cédex, France

E-mail: jean-francois.legall@math.u-psud.fr